arXiv:1510.04721v1 [math.PR] 15 Oct 2015

SITE RECURRENCE FOR COALESCING RANDOM WALK

ITAI BENJAMINI, ERIC FOXALL, ORI GUREL-GUREVICH, MATTHEW JUNGE, AND HARRY KESTEN

Abstract. Begin continuous time random walks from every vertex of a graph and have particles coalesce when they collide. We use a duality relation with the voter model to prove the process is site recurrent on bounded degree graphs, and for Galton-Watson trees whose offspring distribution has exponential tail. We prove bounds on the occupation probability of a site, as well as a general 0-1 law. Similar conclusions hold for a coalescing process on trees where particles do not backtrack.

1. Introduction

Coalescing random walk (CRW) starts with one particle at each vertex of an undirected graph. Each then performs a continuous time nearest neighbor random walk, jumping according to a mean 1 exponential clock. When particles collide they bind together and proceed as one. Say that CRW is *site recurrent* if every site is almost surely visited infinitely often. If instead this occurs with probability 0, call the process *transient*. Our main tool for proving site recurrence is the following necessary and sufficient condition in terms of the expected occupation time of a vertex.

Proposition 1. *Site recurrence is equivalent to infinite expected occupation time at any vertex. Moreover, CRW is either site recurrent or transient (i.e. it satisfies a 0-1 law).*

Let $p_t(v)$ be the probability a particle is at the vertex v at time t, so that site recurrence is equivalent to divergence of $\int p_t(v)$. We use duality with the voter model to obtain non-integrable lower bounds on the following graphs:

Theorem 2. *CRW is site recurrent on:*

(i) Bounded degree graphs. If the maximum degree is D, then for all vertices v

$$p_t(v) \geq (1 + Dt)^{-1}$$

for all $t \geq 0$.

(ii) Galton-Watson trees with offspring distribution on $\mathbb{Z}^+$ and the probability of k offspring bounded by e^{-ck} for some $c > 0$ and large enough k. Here

$$p_t(v) \geq C(t \log t)^{-1}$$

for all vertices v, large enough t, and some $C > 0$ that depends on c.

Note that there are unbounded degree graphs for which CRW is not site recurrent; even the non-coalescing system of independent random walks is transient on trees with rapidly increasing degree. We are not sure how much the exponential tail hypothesis in (ii) can be weakened. See Further Questions (a) for more discussion. A corollary to Proposition 1 is a general upper bound on the probability that a vertex v is unoccupied on the interval (t, u).

Corollary 3. *Let σ_t be the first time after t that v is occupied by a particle. It holds that*

$$\mathbf{P}(\sigma_t > u) \leq \frac{t}{t + \int_t^u p_s(v) ds}.$$

And, for a graph with maximum degree D

$$\mathbf{P}(\sigma_t > u) \leq \frac{t}{t + \frac{1}{D}(\log(1+Du) - \log(1+Dt))} = O(1/\log(u)).$$

We also give a universal upper bound for $p_t(v)$ on general graphs. It follows that the occupation probability decays to 0 for any graph. The upper bound is a small modification of an argument in [Gri78], so we also credit David Griffeath.

Proposition 4 (Griffeath)**.** *Let G be a connected, infinite graph. For all vertices v, any $\epsilon > 0$, and large enough t (depending on ϵ) it holds that $p_t(v) \leq (1+\epsilon)/(2\sqrt{\pi t})$.*

History. The study of coalescing systems began in the 1970's with the paper of Erdős and Ney [EN74]. The duality relationship to the voter model, which we rely heavily upon, was first observed in [HL75]. Variations of coalescing random walk continue to find new applications. For example, random coalescence involving multiple types of particles, and particle interaction rules, is used to model certain chemical reactions (see [Hol83], [BL88] and [vdBK00]). Non-spatial models such as Kingman coalescence ([Kin82]) find applications in modeling ancestry in biology. A survey of coalescence models can be found in [Ber09]. Arratia [Arr83] looks at site recurrence for discrete time walks, and annihilating systems, both with possibly vacant sites in the starting configuration. CRW is applied to the voter model in [BL15]. Also, it is studied in more generality in [RV15] and [GPTZ15]. Other recent articles have focussed on different settings. Its behavior on finite graphs is of interest to computer scientists. The model on the d-dimensional torus is introduced in [Cox89]. There, they study the expected time for the process to coalesce into a single particle. In [CEOR12] the coalescence time is studied on a variety of finite graphs. Elsewhere, in a continuous spatial setting, recurrence is studied with coalescing diffusions by Cabezas, Rolla and Sidoravicious in [CRS13].

Early results for coalescing random walk focused on the lattices $\mathbb{Z}^d$. In [Gri78] Griffeath shows that both coalescing and annihilating random walk on $\mathbb{Z}^d$ is a.s. *weakly recurrent*, under certain restrictions on the vacant sites in the starting configuration. Weakly recurrent means that each site is occupied infinitely often, but for a decreasing fraction of time. An important ingredient in the proof of recurrence is an estimate for the function p_t, the probability a particle occupies the origin at time t. In [BG80], Bramson and Griffeath study p_t in the coalescing case and compute its asymptotics for every $d \geq 1$. Rather nicely, for $d \geq 3$ it holds that $p_t \sim (\gamma_d t)^{-1}$ with γ_d the probability a random walk on $\mathbb{Z}^d$ never returns to its starting point. The proof of this is computational; later, [vdBK02] Kesten gives a probabilistic argument that revolves around the heuristic $p_t' \approx -\gamma_d p_t^2$.

Overview. The main idea is to obtain information about $p_t(v)$ from a dual voter model. This dual was first applied to CRW in [HL75], and subsequently utilized in [HS79, Gri78, Arr81, Arr83]. In Corollary 6 we deduce that $p_t(v)$ is equivalent to the probability a time-changed nearest neighbor simple random walk avoids 0 up to time t. All of our estimates come from studying this random walk.

Further Questions. We record several questions regarding coalescing and annihilating random walk here:

(a) Can the assumptions on the degree in Theorem 2 (ii) be weakened? We expect that our approach extends (at least) to stationary graphs with finite expected degree.
(b) Suppose G is an infinite unimodular random graph in which each vertex is assigned an infinite trajectory in an ergodic invariant way (see [BC12]). Particles, one from each vertex, start moving along their trajectory in continuous time and annihilate when meeting. Is the resulting process recurrent? *Start by showing it on Euclidean lattices.*
(c) Place ϵ-balls (meteors) in Euclidean space with centers according to a unit intensity Poisson process. At time 0 each chooses a direction uniformly randomly and proceeds along this direction at unit speed (non-random). When two meteors collide, they annihilate. Is the origin a.s. occupied by infinitely many meteors for all $d \geq 1$? This is discussed in more detail in Section 3.
(d) Have particles perform annihilating random walk on a graph where the particle started at x steps according to an exponential clock with mean an independent uniform$(0, 1)$ random variable. Is this model on $\mathbb{Z}^d$ still recurrent? If so, what can be said of the limiting speed of the particles visiting the origin? Possibly slower moving particles survive longer, and the average speed of particles visiting the origin decays with time.

Outline. Section 2 starts with the proof of Proposition 1. We also establish, in Lemma 5, that infinite expected occupation time is equivalent to survival of a nearest neighbor random walk in the voter model. Corollary 6 relates this back to p_t. Sections 2.1 and 2.2 contain the proofs of Theorem 2 (i) and (ii), respectively. Section 3 discusses some non-backtracking variants, and contains the proof of site recurrence for a non-backtracking model on bounded degree trees and Galton-Watson trees with exponential tail.

2. Site recurrence for coalescing random walk

Coalescing random walk on a graph, $G = (V, E)$, has a graphical representation as follows: each edge is replaced with two directed edges and an independent Poisson process with unit intensity is placed on each directed edge, indexed by time. When the bell of a Poisson process for the edge (u, w) rings we check if there's a particle at u and if so, we move it to w. If there's already a particle at w, they merge. We denote the process $(\xi_t)_{t\geq 0}$ with $\xi_t \in \{0, 1\}^V$ equal to the set of occupied vertices at time t, and occasionally ξ_t^v for the location at time t of the particle that began at v. In this notation we have $p_t(v) = \mathbf{P}(\xi_t(v) = 1)$ is the probability that v is occupied by a particle at time t. Thus, $\int_0^\infty p_t(v)dt$ is the expected occupation time of v.

Proof of Proposition 1. If there is positive probability of infinite occupation time at v, then the expected occupation time is infinite. For the other direction we generalize [Arr83, Lemma 2].

Suppose that $\int_0^\infty p_t(v) = \infty$. For any $t \in [0, \infty)$ let $\sigma_t = \inf\{s \geq t \colon \xi_t(v) = 1\} \in [0, \infty]$, the first time after t that v is occupied by a particle. We wish to establish that

$$\mathbf{P}(\sigma_t < \infty) = 1, \qquad \forall t \geq 0.$$

The basic coupling $\xi_t^A = \{\xi_t^x \colon x \in A\}$ for $A \subseteq V$ has the property that $A \subseteq B \subseteq V$ implies $\xi_t^A \subseteq \xi_t^B \subseteq \xi_t$, so the Markov process $(\xi_t^A \colon A \subseteq V)$, with state space $\{0, 1\}^V = \{A \colon A \subseteq V\}$ ordered by set inclusion, is attractive. Define $p_t^A = \mathbf{P}(\xi_t^A(v) = 1)$, and also

$I(t,u) = \int_t^u p_s(v)ds$. By assumption,

$$\lim_{u\to\infty} I(t,u) = \infty, \qquad \forall t \geq 0. \tag{1}$$

Let f_{σ_t} be the density function of σ_t and $\mathbf{E}_{\sigma_t}$ denote the expectation taken over all possible realizations of ξ_{σ_t} given that $\sigma_t = r$. Using the strong Markov property we have for $t < u$,

$$\begin{aligned} I(t,u) := \mathbf{E}\int_t^u \mathbf{1}\{\xi_s(v) = 1\}ds &= \int_{r=t}^u f_{\sigma_t}(r)\mathbf{E}_{\sigma_t}\left(\mathbf{E}\int_0^{u-r} \mathbf{1}\{\xi_s^A(v) = 1\}ds\right)dr \\ &\leq \int_{r=t}^u f_{\sigma_t}(r)\left(\mathbf{E}\int_0^u \mathbf{1}\{\xi_s(v) = 1\}ds\right)dr \\ &\leq \mathbf{P}(\sigma_t \leq u)(t + I(t,u)). \end{aligned}$$

Dividing both sides by $t + I(t,u)$ we arrive at the inequality

$$\mathbf{P}(\sigma_t \leq u) \geq I(t,u)/(t + I(0,u)). \tag{2}$$

For fixed t, taking $u \to \infty$ yields $\mathbf{P}(\sigma_t < \infty) \geq 1$ by (1).

We conclude by describing the $0-1$ law. The above argument establishes that if CRW occupies a site for infinite time with positive probability, then it does so with probability 1. As G is assumed to be connected, it follows that all sites are occupied infinitely often with probability 1. Therefore, the process is either site recurrent (recall this is defined as an almost sure event) or transient. □

Proof of Corollary 3. The lower bound on $\mathbf{P}(\sigma_t \leq u)$ at (2) yields an upper bound on the probability v is unoccupied from time t to u:

$$\mathbf{P}(\sigma_t > u) \leq 1 - \frac{I(t,u)}{t + I(t,u)} = \frac{t}{t + \int_t^u p_s(v)ds}. \tag{3}$$

Which is the first part of the corollary. The second part follows by applying the bound on $p_t(v)$ in Theorem 2 (i) and integrating. □

Theorem 2 allows for site recurrence to be deduced by proving $p_t(v)$ is non-integrable. Our approach is to express $p_t(v)$ in another way. Consider the dual process to this model, which is called the *voter model*. In the dual model we start with a partition of the space into clusters, where initially each vertex corresponds to a different cluster. When the bell at (u,w) rings the vertex w is added to the cluster containing u. We denote the process $(\zeta_t^v)_{t\geq 0}$ where ζ_t^v is the set of vertices belonging to the cluster that initially consists of the vertex v. If we run this model in reverse time, from time t to 0, we see the cluster that began at v at time t at time 0 consists of exactly the particles that in the coalescing model are at v at time t. In particular,

$$p_t(v) = \mathbf{P}(\xi_t(v) = 1) = \mathbf{P}(\zeta_t^v \neq \varnothing). \tag{4}$$

The advantage of working with the voter model is that the size of ζ_t^v is a nearest-neighbor symmetric random walk with transition rate depending on the number of boundary edges. Indeed, at any moment there are some directed edges going out of the cluster, and the same number of edges coming in. More precisely, the cluster size is a skip-free process on the integers that moves with rate equal to the size of the current boundary of the cluster. We record this fact in the following lemma. Let $|\cdot|$ denote either the counting measure of a finite set.

Lemma 5. *Define $\zeta_t^v \subseteq V$ to be the set of vertices in the cluster of v at time t. For each $x \in \zeta_t^v$ let $\partial_t(x) = \{(x,y) \in E \colon y \notin \zeta_t^v\}$. Let $\sum_{x\in\zeta_t^v} |\partial_t(x)|$ be the number of edges leading out of ζ_t^v. The process has the following properties:*

(i) $|\zeta_0^v| = 1$.
(ii) *Let $\tau = \inf\{t \colon |\zeta_t| = 0\}$. For all $t \geq \tau$ it holds that $|\zeta_t^v| = 0$.*
(iii) *The process is a martingale that transitions to $|\zeta_t^v| \pm 1$ at rate $\sum_{x\in\zeta_t^v} |\partial_t(x)|$.*

Proof. Properties (i) and (ii) follow from the construction of ζ_t^v. We turn our attention to property (iii). For each $x \in \zeta_t^v$, x is removed from ζ_t^v at rate $|\partial_t(x)|$ and each of the $|\partial_t(x)|$ sites in $\partial_t(x)$ is added to ζ_t^v at rate 1. Since the rates balance,

$$\mathbf{E}[|\zeta_t^v| \mid |\zeta_{t-}^v|] = |\zeta_{t-}^v|,$$

which establishes $|\zeta_t^v|$ is a martingale. Summing the rates over $x \in \zeta_t^v$ shows that $|\zeta_t^v|$ transitions to $|\zeta_t| + 1$, and to $|\zeta_t^v| - 1$, each at rate $\sum_{x\in\zeta_t^v} |\partial_t(x)|$. □

This lets us describe p_t in terms of the voter model.

Corollary 6. $p_t(v) = \mathbf{P}(|\zeta_t(v)| > 0)$. *This is the probability a nearest neighbor random walk with transition rate $\sum_{x\in\zeta_t} |\partial_t(x)|$ and absorbing state at 0 is yet to reach zero at time t.*

It follows that $p_t \to 0$ on any infinite, connected graph.

Proof of Proposition 4. The transition rate in ζ_t^v is always at least two. By Corollary 6 and a straightforward coupling we have p_t is at most $\tilde{p}_t = \mathbf{P}(X_s > 0,\ \forall s \leq t)$, with X_s a rate-2 continuous time simple symmetric random walk started at 1. Using the reflection principle together with the local central limit theorem, $\tilde{p}_t \sim 1/(2\sqrt{\pi t})$ as $t \to \infty$, and the result follows. □

Remark 7. For coalescing walk on $\mathbb{Z}$ with nearest neighbour connections, since ζ_t^v is always of the form $\{x, x+1, ..., x+k\}$ for some $x \in \mathbb{Z}, k \in \mathbb{Z}^+$, its transition rate is exactly $2 \cdot 2 = 4$, so the above inequality is an equality, and gives the exact asymptotics $p_t \sim 1/(2\sqrt{\pi t})$, as observed in [BG80]. Compared to [BG80] there is an extra factor of $1/2$; our convention differs from theirs in that the transition of a particle at v is equal to $\deg v$ and not 1, since in our case $\deg v$ is allowed to vary.

2.1. **Site recurrence for bounded degree graphs.** Now we turn our attention to proving site recurrence on general graphs. Define $\tau_v = \inf\{t \colon \zeta_t^v = \varnothing\}$. Integrating over t in the duality relation (4) we find

$$\mathbf{E}\tau_v = \int_0^\infty \mathbf{P}(|\zeta_t^v| > 0)dt = \int_0^\infty \mathbf{P}(\xi_t(v) = 1)dt = \int_0^\infty p_t(v)dt.$$

So, proving site recurrence is equivalent to showing that the first hitting time of 0 for the simple random walk $|\zeta_t^v|$ (with random and time-varying transition rate) has infinite expectation. We start with the case when G has bounded degree.

Proof of Theorem 2 *(i).* Let $v \in V$ with the maximum degree of vertices in G bounded by D. Lemma 5 establishes that the transition rate of ζ_t^v is less than or equal to $D|\zeta_t^v|$. Let W_t be a continuous time nearest-neighbour random walk on $\mathbb{Z}^+ \cup \{0\}$ with $W_0 = 1$. The walk transitions from $k \in \mathbb{Z}^+ \cup \{0\}$ to $k \pm 1$ each at rate Dk, and is absorbed at 0. Letting $\theta(t) = \mathbf{P}(W_t > 0)$, if follows from Corollary 6 and a straightforward coupling of $|\zeta_t^v|$ with W_t that $p_t(v) \geq \theta(t)$, so it suffices to control $\theta(t)$.

The process W_t can be interpreted as the number of particles in a branching process in which each particle independently dies, or gives birth to a single offspring, each at rate D. Let $\rho(t) = \mathbf{P}(W_t = 0)$. By independence, for each $k \geq 0$ we have

$$\mathbf{P}(W_{t+h} = 0 \mid W_h = k) = \rho(t)^k. \tag{5}$$

Recalling that $W_0 = 1$ then conditioning on W_h for small $h > 0$,

$$\rho(t+h) = \mathbf{P}(W_{t+h} = 0) = \sum_{k \geq 0} \mathbf{P}(W_{t+h} = 0 \mid W_h = k)\mathbf{P}(W_h = k).$$

The event that two or more transitions happens on the interval $[0, h]$ is contained in the event that a rate $2D$ exponential clock rings, then a rate $4D$ exponential clock rings (i.e. we go from 1 to 2 particles then another transition happens). The probability of this is bounded by the probability that $X + Y < h$ for X and Y rate $4D$ exponential random variables, and has density $f_{X+Y}(t) = \lambda^2 t e^{-\lambda t}$ with $\lambda = 1/4D$. Integrating on $[0, h]$, then taking the Taylor expansion we have $\mathbf{P}(X + Y < h) = O(h^2)$. Since we will be dividing by h and letting it tend to 0, we can combine all of the events that occur with two or more transitions into an $O(h^2)$ term. Using the expression (5) this lets us write $\rho(t + h)$ as

$$\rho(t+h) = \underbrace{1 \cdot Dh}_{\text{dies out}} + \underbrace{\rho(t)(1 - 2Dh)}_{\text{no change}} + \underbrace{\rho(t)^2 Dh}_{\text{increases by 1}} + O(h^2).$$

Subtracting $\rho(t)$, dividing by h and taking $h \downarrow 0$ this converges to the equation $\rho' = D(1-\rho)^2$. So, for the survival probability $\theta(t) = 1 - \rho(t)$ we find $\theta' = -\theta^2$, with $\theta(0) = 1$, whose unique solution is $\theta(t) = 1/(1 + Dt)$. □

2.2. **Site recurrence for Galton-Watson trees.** A more general upper bound on the transition rate is

$$|\text{maximum exposed degree}| \cdot |\zeta_t|.$$

Our hypothesis that the offspring distribution of our Galton-Watson tree has exponential tail guarantees that the maximum exposed degree is asymptotically bounded by log(number of steps). Ultimately this lets us compare with the divergent integral $\int_{t_0}^{\infty} (t \log t)^{-1} dt$. This is made rigorous below.

Proof of Theorem 2 *(ii).* Again, by Corollary 6 it suffices to prove that ζ_t^v has infinite expected survival time. For convenience we denote ζ_t^v by ζ_t. Let $H_t = \cup_{s \leq t} \zeta_t$ be the vertices visited up to time t. Define the random times $0 = t_1 < t_2 < \ldots$ as when a vertex is added to H_t, and list them as $v_1, v_2, \ldots$ in the order they are discovered, with v_1 being the root (ρ). The transition rate of $|\zeta_t| \to |\zeta_t| \pm 1$ is at most $M_t|\zeta_t|$ where

$$M_t = \sup\{\deg v \colon v \in H_t\}.$$

So, our first goal is to construct the voter model in such a way that M_t can be easily controlled. A simple way to do this is to construct G "on the fly." That is, let $(X_i)_{i \geq 1}$ be an i.i.d. sequence of copies of the offspring distribution, and at time t_i, sample X_i to determine the offspring distribution of v_i, which is then fixed for all $t > t_i$. This does not disturb the sample path distribution of ζ_t, and has the advantage that the quantity

$$D_k = \max_{i \leq k} \deg v_i$$

is equal to $\max_{i \leq k} X_i$ where X_i is a *fixed* (as opposed to being a randomly indexed) i.i.d. sequence. Since, by assumption, $\mathbf{P}(X_i > x) \leq e^{-cx}$ for some $c > 0$ and large enough x, a

union bound gives $\mathbf{P}(D_k > x) \le ke^{-cx}$, and setting $x = (3/c)\log k$,

$$\mathbf{P}(D_k > (3/c)\log k) \le k^{-2} \tag{6}$$

for large enough k.

Now, let $0 = t_0 < t_1 < t_2 < ...$ denote the jump times of ζ_t. In what follows we will want the set of jump times to be infinite, so if $\zeta_{t_i} = 0$ (i.e. the cluster dies out), just include jumps back to 0 at rate 1. Since $|\xi_i|$ is a martingale with $\mathbf{E}|\zeta_{t_i}| = |\zeta_0| = 1$ for all i, Doob's martingale inequality implies that for all $n > 0$

$$\mathbf{P}(\sup_i |\zeta_{t_i}| > bn) \le (bn)^{-1} \tag{7}$$

for any $b > 0$. Clearly M_{t_i} is nondecreasing in i and $M_{t_i} \le D_{1+i}$, since vertices are exposed one at a time. Thus, the transition rate of $|\zeta_{t_i}|$ is at most $D_{i+1}|\zeta_{t_i}|$. Combining these observations with (6) and (7), we find that with probability at least $1 - (bn)^{-1} - n^{-2}$, for $t \le t_{n^2}$ the transition rate in ζ_t is at most

$$(3/c)\log(n^2+1)bn \le (6b/c)n\log(n+1). \tag{8}$$

Let $m_n = m_n(b,c) = (6b/c)n\log(n+1)$. A quick summary: with high probability the first n^2 transitions happen at rate no more than m_n. Equivalently, the time t_{n^2} is bounded below by the sum of n^2 independent exponentials with rate m_n. This is an Erlang distribution, X_{n^2}, with shape parameter n^2 and rate m_n. Thus, we have

$$\text{mean: } \mu_n = \frac{n^2}{m_n}, \text{ and variance: } \sigma_n = \frac{\mu_n}{m_n}.$$

Chebyshev's inequality guarantees that

$$\mathbf{P}(|X_n - \mu_n| \ge \mu_n/2) \le \frac{\mu_n/m_n}{(\mu_n/2)^2} = \frac{4}{\mu_n m_n} = \frac{4}{n^2}.$$

One side of the above estimate is

$$\mathbf{P}(X_{n^2} \le \mu_n/2) \le 4n^{-2}. \tag{9}$$

By comparison, and using (8) and (9), we have

$$\mathbf{P}(t_{n^2} \ge n^2/(2m_n)) \ge 1 - (bn)^{-1} - n^{-2} - 4n^{-2}.$$

From the well-known first passage distribution for random walk, for the random variable $N = \inf\{n\colon |\zeta_n| = 0\}$ we have $P(N > n^2) \ge c/n$ for some possibly smaller $c > 0$. Note that although m_n depends on c, shrinking c does not affect the estimate. Letting $\tau = \inf\{t\colon |\zeta_t| = 0\}$ as before, and letting $a_n = \mu_n/2 = cn/(12b\log(n+1))$, for $a_{n-1} < t < a_n$,

$$t\log t \ge \frac{c(n-1)}{12b\log n}\left(\log(n-1) - \log(12b/c) - \log\log n\right) \ge \frac{cn}{24b}$$

i.e., $n \le (24b/c)t\log t$ for n large enough. For the same t, then,

$$\mathbf{P}(\tau > t) = \mathbf{P}(|\zeta_t| > 0) \ge \mathbf{P}(|\zeta_{t_{n^2}}| > 0, t_{n^2} > a_n).$$

Since the survival time of the cluster is independent of the rate at which it jumps we have

$$\mathbf{P}(\tau > t) \ge c/n(1 - (bn)^{-1} - O(n^{-2})), \qquad a_{n-1} < t < a_n.$$

The right side is at least $c/(2n)$ for n greater than some n_0. Letting $t_0 = a_{n_0}$ and using the upper bound on n,

$$\mathbf{E}\tau = \int_0^\infty \mathbf{P}(\tau > t)dt \ge \frac{c^2}{48b}\int_{t_0}^\infty \frac{1}{t\log t}dt = \infty.$$

□

3. Non-backtracking coalescing random walk on trees

We are also interested in understanding similar, but less random processes. A lack of symmetry in these settings makes it difficult to apply known techniques. We are hopeful that progress will lead to new ideas.

The *non-backtracking coalescing random walk* is defined in the same way as the coalescing random walk with particles instead performing non-backtracking random walk. More precisely, the state of a particle is specified by a vertex-edge pair $(u, \{u, v\})$, and when an edge clock rings at a directed edge (u, w), the particle moves to w if and only if $w \neq v$. If the particle moves from u to w, its state is updated to $(w, \{w, u\})$, so that its next jump cannot be back to u. It will be convenient to assume that each particle is initialized with a uniformly chosen edge along which it cannot move, that is, the particle initially at v has state $(v, \{v, u\})$ where u is a uniform random neighbour of v. With particles coalescing there is ambiguity about whose path to remember. There are several well-defined ways to assign priority. On a rooted tree, we analyze the special case where we always remember the path of particles moving towards the root. With the model defined in this way we do not quite have a voter model dual, but a closely related process does. Analogous to Proposition 1 we prove a necessary and sufficient condition for site recurrence.

Proposition 8. *Consider coalescing non-backtracking random walk on a rooted tree with priority given to particles moving towards the root. The process is site recurrent at the root if and only if the expected survival time of the cluster in a certain voter model is infinite.*

We can deduce site recurrence on bounded degree trees and some trees with unbounded degree.

Theorem 9. *The process from Proposition* 8 *is site recurrent at the root of either a bounded degree tree or a Galton-Watson tree whose offspring distribution is as in Theorem* 2 *(ii).*

Non-backtracking removes a vital symmetry from the argument. The proof goes by, once again, constructing a dual voter process and showing the cluster of the root survives for an infinite expected amount of time. Our "priority to the root" rule is hand-picked to preserve monotonicity and the existence of a dual voter model. Neither property exists in other equally natural non-backtracking models. Further progress in these different settings will likely require a new approach. Consider the following conjecture:

Conjecture 10. *Non-backtracking coalescing random walk with any priority scheme is site recurrent on bounded degree trees.*

The inspiration for studying non-backtracking processes comes from the following *meteor model* on $\mathbb{R}^d$. Place ϵ-balls in Euclidean space with centers according to a unit intensity Poisson process. At time 0 each chooses a direction uniformly randomly and proceeds along this direction at unit speed (non-random). When two meteors collide, they annihilate.

Conjecture 11. *The origin a.s. is occupied by infinitely many meteors for all* $d \geq 1$.

This problem appears quite difficult. It could be discretized to an annihilating system of random walks by uniformly assigning each particle a geodesics to ∞ from which it never deviates and steps along according to a Poisson clock. The integer lattice is a natural graph to start with. Or, perhaps hyperbolic space – in which random walk paths stay within a logarithmic neighborhood of a geodesic – would be a more tractable place to study this problem.

3.1. **Recovering a dual.** Since the graph is a tree, a particle either moves towards the root, or away, at each jump. Once it has moved away for the first time, at subsequent jumps it must always move away, since the only way back towards the root requires backtracking. To simplify matters, we suppose that at each vertex, the initial forbidden edge is chosen uniformly from the edges that lead away from the root.

Since coalescing particles may have different histories we must decide which one to remember. Therefore, upon collision we define the following rule for annihilating exactly one of the two colliding particles; in the original setting with memoryless walks, any such rule leads to the coalescing model.

- If one particle is moving towards the root, and the other particle is moving away, annihilate the particle moving away.
- Otherwise, annihilate the particle currently occupying the site (i.e. keep the particle that just moved).

As currently stated, this process depends on past information. Running the process in reverse would require information about the future. Thus, it does not have a dual voter model. Still, a simple observation yields a related model that does have a dual.

Let X_T be the occupation time of the root up to time T. Notice that particles moving away from the root are inert; based on the rule above they cannot block upward moving particles, and as noted before they cannot revisit the root. Thus, X_T is unchanged if we delete particles the instant they turn away from the root. This modification gives the following model; for a vertex v, let $d_v = \deg v$.

- Suppose v is not the root.
  - A particle at v moves towards the root at rate 1.
  - A particle at v is deleted at rate $d_v - 2$.
- Suppose $v = \rho$ is the root. A particle at ρ is deleted at rate $d_\rho - 1$.

We can think of this model as follows: each particle attempts to travel up to the root, coalescing with other particles upon collision, and particles (or coalesced collections of particles) are instantaneously zapped out of existence at some rate that depends on their present location.

We can simplify the description somewhat by introducing a single absorbing vertex $\mathfrak{a}$ and considering the process on $V \cup \{\mathfrak{a}\}$, where every vertex has a directed edge pointing to v which rings at rate described below.

Particles at $\mathfrak{a}$ do not move. The transitions for particles at $v \in V$ are as follows.

- Move towards the root at rate 1.
- Move directly to $\mathfrak{a}$ at rate $d_v - 2$.

A graphical representation of this model can be obtained by placing an independent Poisson process with rate 1 at each upward directed edge, and with rate $d_v - 2$ at each vertex v.

The model enjoys the same monotonicity as the coalescing random walks – resetting to the initially full configuration maximizes the probability of occupying the root in the future. Then, as for the coalescing random walks, there is a dual voter model. In this case, deletion of a particle at v corresponds to the addition of v to the cluster of $\mathfrak{a}$. Note that since the direction of motion is reversed in the voter model, clusters on the tree must expand away from the root. Altogether, the voter model has the following transitions. A *down-going* directed edge (w, u) is an edge directed away from $\mathfrak{a}$. These are the rules we use in the proof of Lemma 12. We box it for emphasis:

- Along each down-going edge (w, u), at rate 1, u is added to the cluster containing w.
- At each vertex v on the tree, at rate $d_v - 2$, v is added to the cluster of $\mathfrak{a}$.

The existence of a dual lets us establish Proposition 8.

Proof of Proposition 8. We have established that X_T in the non-backtracking coalescing random walk has the same distribution as in the simpler model. Since the simpler model has monotonicity and a voter model dual, the same argument used to prove Proposition 1 gives the desired equivalence. □

3.2. **Site recurrence.** Now we can turn our attention to proving infinite expected survival time of the root cluster in the voter model. The voter model from Section 3.1 also has the martingale property. As before, let ζ_t^v denote the cluster that began at v.

Lemma 12. *Consider the voter model $\zeta_t^v \in V \cup \{\mathfrak{a}\}$ described above. For each $v \in V$, so long as $|\zeta_t^v| < \infty$, the size of the cluster ζ_t^v is a martingale. It transitions to $|\zeta_t| \pm 1$ at rate $\sum_{w\in\zeta_t}(d_w - 1) - 2|\{(w, y)\colon w, u \in \zeta_t^v\}|$.*

Proof. Fix a vertex $v \in V$ for which we will consider the cluster $\zeta_t = \zeta_t^v$. Let $r_t^+ = r_t^+(v)$ and $r_t^- = r_t^-(v)$ denote the rate at which $|\zeta_t| \to |\zeta_t| \pm 1$, respectively. Note that the transition rules prohibit $\mathfrak{a} \in \zeta_t$. Moreover, ζ_t is unchanged if we assume that initially, all vertices but v belong to the cluster of $\mathfrak{a}$. Therefore, it is enough to check that for any finite W, if $\zeta_t = W$ and $\zeta_t^{\mathfrak{a}} = V \cup \{a\} \setminus W$ then $r_t^+ - r_t^- = 0$. For a vertex $w \neq \mathfrak{a}$ let $\hat{w}$ denote its unique parent vertex, i.e., the unique vertex such that there is a down-going edge to w, and let $\mathfrak{o}(w)$ denote the set of childs vertices. From the transition rules it follows that

$$r_t^+ = \sum_{w\in W}(d_w - 1) - |\{u \in \mathfrak{o}(w)\colon u \in W\}|$$

and

$$r_t^- = \sum_{w\in W}(d_w - 2) + \mathbf{1}(\hat{w} \notin W) = \sum_{w\in W}(d_w - 1) - \mathbf{1}(\hat{w} \in W)$$

and since $\sum_{w\in W} |\{u \in \mathfrak{o}(w) : u \in W\}|$ and $\sum_{w\in W} \mathbf{1}(\hat{w} \in W)$ are both equal to $|\{(w, u)\colon u, w \in W\}|$,

$$r_t^+ - r_t^- = -\sum_{w\in W} |\{u \in \mathfrak{o}(w)\colon u \in W\}| + \sum_{w\in W} \mathbf{1}(\hat{w} \in W) = 0.$$

□

Proof of Theorem 9. With Lemma 12 we can bound the transition rate of $|\zeta_t^\rho|$ by $\sum_{v\in\zeta_t^\rho} d_v$. To prove the part of Theorem 9 concerning bounded degree trees we can follow the same approach as Theorem 2 (i); again we use the transition rate bound $D \cdot |\zeta_t^\rho|$. Similarly, we can use the same technique as the proof of Theorem 2 (ii) to deduce site recurrence for Galton-Watson trees. □

References

[Arr81] Richard Arratia, *Limiting point processes for rescalings of coalescing and annihilating random walks on z^d*, Ann. Probab. **9** (1981), no. 6, 909–936.

[Arr83] ———, *Site recurrence for annihilating random walks on $\mathbb{Z}^d$*, Ann. Probab. **11** (1983), no. 3, 706–713.

[BC12] Itai Benjamini and Nicolas Curien, *Ergodic theory on stationary random graphs*, Electron. J. Probab. **17** (2012), no. 93, 1–20.

[Ber09] Nathanaël Berestycki, *Recent progress in coalescent theory*, Ensaios Matematicos **16** (2009), no. 1.

[BG80] Maury Bramson and David Griffeath, *Asymptotics for interacting particle systems on* $\mathbb{Z}^d$, Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete **53** (1980), no. 2, 183–196 (English).

[BL88] Maury Bramson and Joel L. Lebowitz, *Asymptotic behavior of densities in diffusion-dominated annihilation reactions*, Phys. Rev. Lett. **61** (1988), 2397–2400.

[BL15] M. Balázs and A. László Nagy, *Dependent augmented Branching Annihilating Random Walk*, ArXiv e-prints (2015).

[CEOR12] Colin Cooper, Robert Elsässer, Hirotaka Ono, and Tomasz Radzik, *Coalescing random walks and voting on graphs*, CoRR **abs/1204.4106** (2012).

[Cox89] J. T. Cox, *Coalescing random walks and voter model consensus times on the torus in* $\mathbb{Z}^d$, Ann. Probab. **17** (1989), no. 4, 1333–1366.

[CRS13] M. Cabezas, L. T. Rolla, and V. Sidoravicius, *Recurrence and Density Decay for Diffusion-Limited Annihilating Systems*, ArXiv e-prints (2013).

[EN74] P. Erdos and P. Ney, *Some problems on random intervals and annihilating particles*, Ann. Probab. **2** (1974), no. 5, 828–839.

[GPTZ15] B. Garrod, M. Poplavskyi, R. Tribe, and O. Zaboronski, *Interacting particle systems on Z as Pfaffian point processes I - annihilating and coalescing random walks*, ArXiv e-prints (2015).

[Gri78] David Griffeath, *Annihilating and coalescing random walks on d*, Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete **46** (1978), no. 1, 55–65 (English).

[HL75] Richard A. Holley and Thomas M. Liggett, *Ergodic theorems for weakly interacting infinite systems and the voter model*, Ann. Probab. **3** (1975), no. 4, 643–663.

[Hol83] Richard Holley, *Two types of mutually annihilating particles*, Advances in Applied Probability **15** (1983), no. 1, 133–148 (English).

[HS79] R. Holley and D. W. Stroock, *Central limit phenomena of various interacting systems*, Annals of Mathematics **110** (1979), no. 2, pp. 333–393 (English).

[Kin82] J.F.C. Kingman, *The coalescent*, Stochastic Processes and their Applications **13** (1982), no. 3, 235 – 248.

[RV15] B. Rath and D. Valesin, *Percolation on the stationary distributions of the voter model*, ArXiv e-prints (2015).

[vdBK00] J. van den Berg and Harry Kesten, *Asymptotic density in a coalescing random walk model*, Ann. Probab. **28** (2000), no. 1, 303–352.

[vdBK02] ———, *Randomly coalescing random walk in dimension* ≥ 3, 1–45 (English).

Department of Mathematics, University of Washington
E-mail address: itai.benjamini@weizmann.ac.il

Department of Mathematical and Statistical Sciences, Arizona State University
E-mail address: Eric.Foxall@asu.edu

Einstein Institute of Mathematics, Hebrew University of Isreal
E-mail address: origurel@math.huji.ac.il

Department of Mathematics, University of Washington
E-mail address: jungem@math.washington.edu

Department of Mathematics, Cornell University
E-mail address: kesten@math.cornell.edu